

Towards an Understanding of the Value of Cooperation in uncertain world

Yonghong Wang*, Katia Sycara* and Paul Scerri*

*The Robotics Institute, Carnegie Mellon University, 5000 Forbes Ave, Pittsburgh, PA 15213-3890

Email:yhwang_y2k@yahoo.com

Abstract—Recent work has provided tantalizing hints that small amounts of cooperation may actually hurt a group’s performance rather than help it. In this paper, we take a systematic look at the value of cooperation. Using a simple cooperative task where agents can act effectively individually but where high levels of cooperation will intuitively lead to better behavior, we investigated when and how cooperation helped overall performance. We systematically varied properties of the environment, e.g., the amount of uncertainty, and how much of what sort of cooperation the agents would perform, e.g., information sharing, resource allocation etc. Our experimental results show that, even if communication were free, and while typically coordination helps the team, under some circumstances, the team may be better off not coordinating at all than coordinating a little bit. We show that the level of uncertainty the agents face and their initial understandings of the environment determine whether a small amount of coordination is useful.

Keywords—cooperation; coordination; information; uncertainty

I. INTRODUCTION

Recently, some work has appeared in the literature suggesting that doing only a little cooperation can actually harm overall performance. Taylor et al [4] showed that distributed constraint optimization with cooperation on value choices between small numbers of agents could actually make the optimization process worse. In other work, Velagapudi et al [6] showed that when agents were cooperatively path planning, there were some messages about obstacles that negatively influenced performance, if sent, despite being accurate. These results show that cooperation might not be helpful at all, let alone worth any costs of performing the cooperation. In this paper, we start an effort toward a better understanding of exactly when cooperation is useful and when it is not.

To investigate the value of cooperation, we developed a simulation model where agents can act independently but where overall performance will intuitively be better if the agents cooperate. *The key result of extensive experimentation with the simulation is that there are a variety of cases where a small amount of cooperation or shared information can actually decrease overall performance, even when the information was correct.* If the agents were able to cooperate extensively, performance was always better than

acting individually, but it requires significant amounts of communication.

II. AGENT DECISION AND COOPERATION

Model: We call the task Treasure Hunt in which each agent is assigned an item to acquire. After the item is acquired, the agent will be assigned another item. Items are randomly placed at different locations which are connected by routes. Each item can have multiple instances at different locations. Once an item is collected, it is relocated randomly. If not collected, an item can only stay in its location for certain time, then it is relocated randomly. Agents need to travel from location to location in order to find the items they are assigned. The cost of traveling a route is probabilistic. Agents form a social network which might not depend on their physical locations. Agents can communicate with their neighbors. The goal for each individual agent is to collect as many items as possible and for the team to collect as many items as possible. The agents are completely selfless, having exactly as much desire for the team to collect items as to collect them individually.

In this section we describe the individual decision-making and the cooperation that the agents perform. It is intuitively expected that each type of cooperation helps overall performance.

Agent individual decision: Random: agents pick a nearby location randomly to explore. Path Planning: if the agent has information about the location of the assigned item, it calculates the shortest path to the assigned item using all information available. Since information might change while agent is moving, we use A^* algorithm to generate optimal paths. If the agent has no information about the location of the assigned item, it pick a nearby location randomly to explore or pick one location that has not been explored by other agents.

Ways to Cooperate: Next, we describe the possible ways an agent could cooperate. Notice that a team may employ multiple types of cooperation for better performance.

- Share information about item locations: whenever an agent moves to a new location, he broadcasts what items are at this location to all his neighbors. His neighbors may pass this information to their neighbors until a fixed depth limit is reached.

- **Sharing route costs:** when an agent travels a route, it broadcasts the cost to all its neighbors, and its neighbors may pass the information on to their neighbors until a fixed depth limit is reached.
- **Coordination of exploration:** agents help with each other to explore unknown locations. Whenever an agent is going to explore a new location, he tells all his neighbors what he is going to explore and the neighbors may pass this information to all their neighbors until a fixed depth limit is reached. Whenever an agent receives such kind of information, he doesn't need to explore that location and he may chose to explore other locations
- **Global knowledge:** As an upper bound on performance, we created a centralized, all-knowing algorithm. This algorithm approximates perfect, complete cooperation, but without the high cost

Agent's Prior Beliefs: The cost of a route r is a probabilistic function $c_r > 0$ which is the absolute value of a normal distribution with mean u_r , variance σ_r and $u_r = d_r + m_r$, d_r is the physical Euclidean distance of the route, $m_r > 0$ is a mean shift, and it is the absolute value of a randomly generated number from a normal distribution with mean 0 and variance h , where h is a constant. The variance σ_r is randomly generated from 0 to σ where σ is a constant. It turns out that agent's initial beliefs about uncertain route information are important in determining whether some information sharing is helpful. We consider three different models covering a spectrum of assumptions.

- **Optimistic:** optimistic agents assume a lower limit of the cost, which is the physical Euclidean length of the route, $c_r = d_r$
- **Pessimistic:** pessimistic agent assume an upper limit of the cost, $c_r = d(r) + 3h + 3\sigma$
- **Rational:** rational agents assume route cost to be the expected value, $c_r = (d_r + \sqrt{\frac{2}{\pi}h})(1 - \sqrt{\frac{1}{2\pi}} \int_{-\infty}^{-u} e^{-x^2/2} dx) + \sqrt{\frac{2}{\pi}} e^{-u^2/2} \frac{\sigma}{2}$.

III. RESULTS

Treasure Hunt simulator first generates a map consisting of 500 locations, which are randomly placed in a 200×200 area. Each location is connected to an average of 4 other nearby locations. 500 different types of items, 3 items per type, are allocated randomly to those locations. 500 agents form a social network and are randomly deployed at those locations. We assume the broadcasting information and planning altogether takes 1 unit of time.

Broadcasting Depth: When agents broadcast information to their neighbors, their neighbors can pass the information to their neighbors, until a predefined number of hops is reached.




Figure 1. Value of sharing road cost and variation

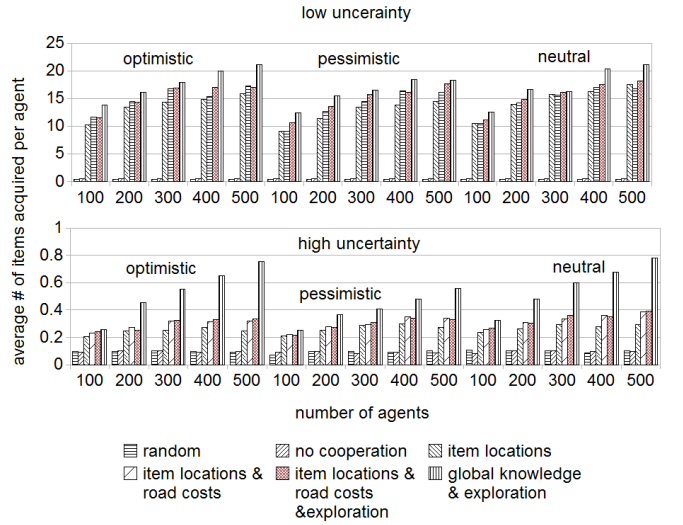


Figure 2. Performance versus number of agents, degree of uncertainty and agent prior assumptions.

We vary the number of hops from 1 to 6 in our experiment.

As Figure 1 shows, the system's performance increases dramatically first, then gradually increases as the number of hops gets higher. When it reaches about 3, or 4, depending on the number of agents, the increase is not very obvious. When the number of agents is low, after 3 hops, the higher hops does not gain very much. When the number of agents is high, 4 hops is good enough to get maximum performance.

Degree of cooperation: In this experiment, we study six different cases: agents move randomly, agents plan individually without communication, agents share information about items at locations, agents share information about items at locations and cost of routes, agents share information about items at locations and cost of routes and exploration of unknown sites, agents share all information with all other agents which is called global knowledge.

Results in Figure 2 show that more cooperation generally




Figure 3. Value of sharing road cost and time

results in better performance. But there are few exceptions. There are cases in which sharing route costs actually degrades performance of neutral agents and there are cases in which coordination of exploration worsens team performance too. Global knowledge improves system performance in all cases, but the value of global knowledge depends on the uncertainty, agent's prior beliefs and the number of agents. It has more value when there are more agents, higher uncertainty.

Cooperation With High Uncertainty: We expect that sharing too few route cost information in a highly uncertain world might hurt team performance. We set up two experiments to verify our hypothesis. First experiment investigates the value of cooperation versus time and the second experiment investigates the value of cooperation versus number of agents.

In this experiment, agents are optimistic. All locations are placed in a square with length 10, shift variation $h = 10$, and variation $\sigma = 200$, which means the uncertainty is high. We divide the game into 10 stages, with 200 units of time each. So in earlier stages, there is little information being shared. As shown in Figure 3, agents without sharing route cost outperform agents sharing road cost in the first two stages, after 3rd stage, agents sharing route cost begin to perform better. Which agrees with our hypothesis.

In this experiment, the number of agents varies from 25 to 500. All locations are in a 20×20 square, shift variation $h = 20$ and $\sigma = 50$, time span is 1000 units. As in Figure 4, when the number of agents is few, the cooperation worsens the system performance, but when the number of agents increases, the value of broadcasting route cost increases. This verifies our hypothesis since when the number of agents is low, the amount of information being shared is low.

Degree of uncertainty: We expect that the value of sharing cost of routes decreases when the uncertainty of the route cost increases. In this experiment, all locations are in a 10×10 square, shift variation $h = 10$ and short game length at 500 units of time. The number of agents is 500. The degree of uncertainty is represented by σ , the variation of




Figure 4. Value of sharing road cost and number of agents




Figure 5. Value of sharing route cost and uncertainty of route cost

route cost. We vary σ from 1 to 200. As in figure 5, when the degree of uncertainty increases, the value of cooperation decreases, and eventually, it will make system performance worse. Sharing route cost suffers more performance drop than not sharing route cost when uncertainty increases, and when the uncertainty reaches a certain level, it is better not to share route cost. As for optimistic agents, sharing route cost outperforms not sharing, until $\sigma = 100$, and after that, sharing route cost make system performance worse. As for pessimistic and neutral agents, sharing route cost stops helping at a very early stage, $\sigma = 30$. In this sense, optimistic agents are more robust against uncertainty, while pessimistic and neutral agents are much more sensitive to uncertainty.

Coordination of Exploring Unknown Locations:




Figure 6. Exploration coordination versus no coordination when the number of agents varies

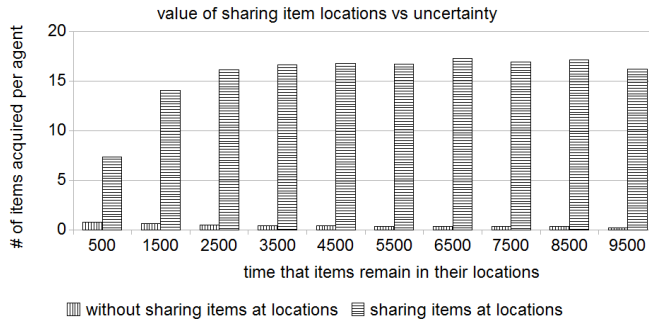


Figure 7. Sharing item at locations versus not sharing when items at location uncertainty varies

Hypothesis: The value of coordination of exploring unknown locations decreases as the number of agents increases. In order to verify our hypothesis, we vary the number of agents from 50 to 1000, and we can see from Figure 6, that the value of coordination gradually decreases when the number of agents increases, until it does not help at all.

Uncertainty of Item Locations:

Hypothesis: The value of broadcasting items at locations decreases when the uncertainty of item locations increases. We define the uncertainty of item locations as the time that an item stays in certain location before it moves to another location. We vary the stay interval from 500 to 10000. The length of the experiment is 10000. As in Figure 7, the value of sharing items at locations increases when the stay interval increases, or the uncertainty of item locations decreases.

IV. RELATED WORK

Teamwork has been widely studied in Distributed AI, [2], [3]. Coordination among cooperative agents is a major challenge in MAS (Multi-Agent System). Much work focuses on defining theoretical foundations to guide agent cooperation and coordination, [5],[6]. It is not realized until recently

that the cooperation may not increase team performance, or even worsen it under certain circumstances. [1] studied the effectiveness of teamwork when the level of cooperativeness or the size of local cooperation group varies. They show that increased cooperation may not always increase performance. Taylor et al [4] also concluded that cooperation among team members may actually degrade the system performance sometimes even the cost associated with the coordination is free.

V. CONCLUSIONS

In this paper, we used a simple game to systematically explore the value of cooperation. As with some previous work, we found cases where a small amount of cooperation actually hurts performance. This is strongly influenced by the prior beliefs of the agents, specifically how their prior beliefs relate to incoming information. This result should make developers think carefully when adding cooperation to a system – the impact may not be as expected. In future work, we plan to analytically show when and how cooperation can be unhelpful.

REFERENCES

- [1] M. Benisch and N. Sadeh. Examining dcsp coordination tradeoffs. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, May 2006.
- [2] B. J. Grosz and S. Kraus. Collaborative plans for complex group action. *Artificial Intelligence*, 86(2):269–357, 1996.
- [3] N.R. Jennings. Commitments and conventions: The foundation of coordination in multi-agent systems. *The Knowledge Engineering Review*, 8(3):223–250, 1993.
- [4] Matthew E. Taylor, Manish Jain, Yanquin Jin, Makoto Yooko, and Milind Tambe. When should there be a “me” in “team”? Distributed multi-agent optimization under uncertainty. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, May 2010.
- [5] Prasanna Velagapudi, Student Member, Katia Sycara, and Paul Scerri. Decentralized prioritized planning in large multirobot teams. *Robotics*, pages 4603–4609, 2010.
- [6] Prasanna Velagapudi, Oleg Prokopyev, Paul Scerri, and Katia Sycara. Analyzing the performance of randomized information sharing. In *Proceedings of 8th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2009)*, May 2009.